1. (currently amended) A method of determining which of a set of content provider mirror sites should receive an end user's initial ~~content~~ request, comprising:

generating a network map during a map generation process according to the following sub-steps:

identifying a set of proxy points, wherein each proxy point represents a given point in the Internet at which a trace originating from each of a set of content provider mirror sites directed toward a given name server intersect;

probing the proxy points to generate given data ~~for each proxy point~~;

generating a download predictor score for each content provider mirror site based on the given data generated by probing the proxy points;

identifying which mirror site provides a best download performance based on the download predictor score; and

associating a given name server IP address with the identified mirror site to generate the network map; and

upon completion of the map generation process, in response to an end user's initial ~~content~~ request to a given local name server, returning an IP address of the identified mirror site based on information in the network map.

2. (currently amended) A method of optimizing a user's initial request to a content provider web site that is replicated at a set of mirror sites, comprising:

responsive to an end user's local name server making a request ~~to~~ associated with the content provider's web site, directing the request to a global load balancing service having a network map that estimates relative connectivity to the mirror sites from ~~each of~~ a set of one or more proxy points, ~~wherein each proxy point represents a set of local name servers~~ wherein each proxy point represents a given point in the Internet at which a trace route over the Internet originating from each of the set of mirror sites directed toward a given name server intersect; and

using the network map to return to the end user's local name server an IP address identifying an optimal mirror site at which the request may be serviced.

3. (cancelled).

4. (currently amended) A method of routing a user's initial request to a content provider web site that is replicated at a set of mirror sites, comprising:

responsive to an end user's local name server making a request ~~to~~ associated with the content provider web site, directing the request to a global load balancing service having a network map that estimates relative connectivity to the mirror sites from ~~each of~~ a set of one or more proxy points, ~~wherein each proxy point represents a set of local name servers~~ wherein each proxy point represents a given point in the Internet at which a trace route over the Internet originating from each of the set of mirror sites directed toward a given name server intersect;

determining whether the network map includes data associating the end user's local name server to one of the mirror sites; and

if not, identifying a given mirror site to respond to the request using a default routing mechanism.

5. (original) The method as described in Claim 4 wherein the default routing mechanism is BGP.

6. (original) The method as described in Claim 4 wherein the default routing mechanism is geo-routing.

7. (currently amended) A method, operated by a server provider, for managing global traffic redirection for a set of content providers operating mirrored sites, wherein the service provider is distinct from the set of content providers operating the mirrored sites, comprising:

for each content provider in the set of content providers, generating a network map during a map generation process according to the following sub-steps:

from each of a set of data centers that host mirrored sites for the content provider,

executing a given network test against each of a set of one or more core points~~, wherein each core point represents a set of name servers~~;

generating a time-weighted average of a given network performance metric based on data generated by executing the given network test;

generating a score for each data center per core point;

generating a set of candidate data centers for each of a set of name servers;

associating a candidate data center to each of a set of IP address space blocks to generate a the network map;

providing the network map to a name server; and

using the network ~~map~~ maps generated for the set of content providers to direct end user requests to a mirrored site to a given data center.

8. (original) The method as described in Claim 7 wherein the given network test is a ping test.

9. (original) The method as described in Claim 7 wherein the given metric is latency or packet loss.

10. (original) The method as described in Claim 7 further including the step of discarding from the set of candidate data centers any data center that does not meet a given operating criteria.

11. (original) The method as described in Claim 10 wherein the given operating criteria is evaluated using a file download test.

12. (currently amended) A method of optimizing a client request to a ~~content provider site~~ domain that is replicated at a set of mirror sites, comprising:

generating a network map that estimates relative connectivity to the mirror sites from ~~each of~~ a set of one or more proxy points~~, wherein each proxy point represents a set of local~~

~~name servers~~ wherein each proxy point represents a given point in the Internet at which a trace route over the Internet originating from each of the set of mirror sites directed toward a given name server intersect;

responsive to a local name server making a request to associated with the ~~content provider's site~~ domain, directing the request to a global load balancing service; and

having the global load balancing service use the network map to return to the local name server an IP address identifying an optimal mirror site at which the request may be serviced.

13. (currently amended) The method as described in Claim 12 wherein the client request originates at a client machine and the ~~content provider site is a~~ domain is associated with a content provider Web site.

14. (original) The method as described in Claim 12 wherein the client request originates from ~~a cache~~ a content delivery network edge server and the set of mirror sites ~~comprises a plurality of~~ comprise storage servers.

15. (original) The method as described in Claim 12 wherein the client request originates at a streaming server and the set of mirror sites comprises a plurality of signal acquisition points.

16. (original) The method as described in Claim 12 wherein the client request originates at a logging process and the set of mirror sites comprises a plurality of log archival servers.

17. (original) The method as described in Claim 12 wherein the client request originates at a mail process and the set of mirror sites comprises a plurality of mail servers.